[EDAS-CFP] VHPC'22 : deadline April 19, rolling abstract : 17th Workshop on Virtualization in HPC and Cloud

Expéditeur : dist@vhpc.org (dist=vhpc.org@edas.info)

À : k.alahassa123@yahoo.fr

Date : vendredi 8 avril 2022 à 03:33 UTC

[You are receiving this call-for-papers from EDAS, based on your user profile. To unsubscribe from call-for-paper messages such as this one, go to `My profile` and update your areas of interest.]

========== CALL FOR PAPERS

17th Workshop on Virtualization in High-Performance Cloud Computing (VHPC 22) held in conjunction with the International Supercomputing Conference - High Performance, May 29 - June 2, 2022. (Springer LNCS Proceedings)

==========

Date: June 2, 2022 Workshop URL: https://vhpc.org

Rolling abstract registration April 19, 2022 - Paper submission deadline (Springer LNCS)

Call for Papers

Containers and virtualization technologies constitute key enabling factors for flexible resource management in modern data centers, and particularly in cloud environments. Cloud providers need to manage complex infrastructures in a seamless fashion to support the highly dynamic and heterogeneous workloads and hosted applications customers deploy. Similarly, HPC environments have been increasingly adopting techniques that enable flexible management of vast computing and networking resources, close to marginal provisioning cost, which is unprecedented in the history of scientific and commercial computing. Most recently, Function as a Service (Faas) and Serverless computing, utilizing lightweight VMs-containers widen the spectrum of applications that can be deployed in a cloud environment, especially in an HPC context. Here, HPC-provided services become accessible to distributed workloads outside of large cluster environments.

Various virtualization-containerization technologies contribute to the overall picture in different ways: machine virtualization, with its capability to enable consolidation of multiple underutilized servers with heterogeneous software and operating systems (OSes), and its capability to live-migrate a fully operating virtual machine (VM) with a very short downtime, enables novel and dynamic ways to manage physical servers; OS-level virtualization (i.e., containerization), with its capability to isolate multiple user-space environments and to allow for their coexistence within the same OS kernel, promises to provide many of the advantages of machine virtualization with high levels of responsiveness and performance; lastly, unikernels provide for many virtualization benefits with a minimized OS/library surface. I/O Virtualization in turn allows physical network interfaces to take traffic from multiple VMs or containers; network virtualization, with its capability to create logical network overlays that are independent of the underlying physical topology is furthermore enabling virtualization of HPC infrastructures.

Publication

Accepted papers will be published in a Springer LNCS proceedings volume.

Topics of Interest

The VHPC program committee solicits original, high-quality submissions related to virtualization across the entire software stack with a special focus on the intersection of HPC, containers-virtualization and the cloud.

Design / Architecture:

- Containers and OS-level virtualization (LXC, Docker, rkt, Singularity, Shifter)
- Hypervisor support for heterogeneous resources (GPUs, co-processors, FPGAs, etc.)
- Hypervisor extensions to mitigate side-channel attacks ([micro-]architectural timing attacks, privilege escalation)
- VM & Container trust and security models
- Multi-environment coupling, system software supporting in-situ analysis with HPC simulation
- Cloud reliability, fault-tolerance and high-availability
- Energy-efficient and power-aware virtualization
- Containers inside VMs with hypervisor isolation
- Virtualization support for emerging memory technologies
- Lightweight/specialized operating systems in conjunction with virtual machines
- Hypervisor support for heterogeneous resources (GPUs, co-processors, FPGAs, etc.)
- Novel unikernels and use cases for virtualized HPC environments
- ARM-based hypervisors, ARM virtualization extensions

Management:

- Container, VM and data management for HPC and cloud environments
- HPC services integration, services to support HPC
- Service and on-demand scheduling & resource management
- Dedicated workload management with VMs or containers
- Workflow coupling with VMs and containers
- Unikernels and lightweight VM application management
- Environments and tools for operating containerized environments (batch, orchestration)
- Novel models for non-HPC workload provisioning on HPC resources

Performance Measurements and Modeling:

- Performance improvements for or driven by unikernels
- Optimizations of virtual machine monitor platforms and hypervisors
- Scalability analysis of VMs and/or containers at large scale
- Performance measurement, modelling and monitoring of virtualized/cloud workloads
- Virtualization in supercomputing environments, HPC clusters, HPC in the cloud
- Energy-efficient deployment of high-performance, ultra-low latency and real-time workloads in cloud infrastructures
- Modeling, control and isolation of end-to-end performance for parallel & distributed cloud/HPC applications

Configuration / Tooling:

- Tool support for unikernels: configuration/build environments, debuggers, profilers
- Job scheduling/control/policy and container placement in virtualized environments
- Measuring and controlling “OS/Virtualization noise”
- Operating MPI in containers/VMs and Unikernels
- GPU virtualization operationalization

In addition to the general research topics above, VHPC'22 encourages particularly contributions on the following focus topics:

- Container Platforms (Kubernetes, Docker, Singularity, Shifter, rkt, …) for Scientific Workflows
- Composable Lightweight Applications and Unikernel Frameworks
- Latency Control and Data/Container Placement in Heterogeneous HPC Virtualized Environments
- Energy-efficiency and Service Orchestration in Virtualized Cloud & HPC Infrastructures

The Workshop on Virtualization in High-Performance Cloud Computing (VHPC) aims to bring together researchers and industrial practitioners facing the challenges posed by virtualization in order to foster discussion, collaboration, mutual exchange of knowledge and experience, enabling research to ultimately provide novel solutions for virtualized computing systems of tomorrow.

The workshop will be one day in length, composed of 20 min paper presentations, each followed by 10 min discussion sections, plus lightning talks that are limited to 5 minutes. Presentations may be accompanied by interactive demonstrations.

Important Dates

April 12, 2022 - Abstract registration April 19, 2022 - Paper submission deadline (extended) (Springer LNCS) May 3, 2022 - Acceptance notification June 2, 2022 - Workshop Day July 10, 2022 - Camera-ready version due

Chair

Michael Alexander (chair), BOKU, Vienna and OeAW, Austria Anastassios Nanos (co-chair), Nubificus Ltd., UK Tommaso Cucinotta, Scuola Superiore Sant'Anna, Italy

Program committee

Stergios Anastasiadis, University of Ioannina, Greece Paolo Bonzini, Red Hat, Italy Eduardo César, Universidad Autonoma de Barcelona, Spain Taylor Childers, Argonne National Laboratory, USA Stephen Crago, USC ISI, USA François Diakhaté CEA DAM Ile de France, France Kyle Hale, Illinois Institute of Technology, USA Brian Kocoloski, University of Southern California, USA Giacomo Lanciano, Scuola Normale Superiore, Italy Giuseppe Lettieri, University of Pisa, Italy Klaus Ma, Huawei Technologies, China Alberto Madonna, Swiss National Supercomputing Center, Switzerland Nikos Parlavantzas, INSA Rennes, France Kevin Pedretti, Sandia National Laboratories, USA Amer Qouneh, Western New England University, USA Carlos Reaño, Universitat de València, Spain Riccardo Rocha, CERN, Europe Lutz Schubert, Ulm University, Germany Borja Sotomayor, University of Chicago, USA Jonathan Sparks, Cray, USA Kurt Tutschku, Blekinge Institute of Technology, Sweden John Walters, University of Southern California, USA Yasuhiro Watashiba, Osaka University, Japan

Paper Submission-Publication

Papers submitted to the workshop will be reviewed by at least two members of the program committee and external reviewers. Submissions should include abstract, keywords, the e-mail address of the corresponding author, and must not exceed 10 pages, including tables and figures at a main font size no smaller than 11 point. Submission of a paper should be regarded as a commitment that, should the paper be accepted, at least one of the authors will register and attend the conference to present the work. Accepted papers will be published in a Springer LNCS volume.

The format must be according to the Springer LNCS Style. Initial submissions are in PDF; authors of accepted papers will be requested to provide source files.

https://edas.info/N29178

Lightning Talks

Lightning Talks are non-paper track, synoptical in nature and are strictly limited to 5 minutes. They can be used to gain early feedback on ongoing research, for demonstrations, to present research results, early research ideas, perspectives and positions of interest to the community. Submit abstract via the main submission link.

General Information

The workshop is one day in length and will be held in conjunction with the International Supercomputing Conference - High Performance (ISC) 2022, May 29 - June 2.